AMENDMENTS TO THE CLAIMS

1. (Currently amended): A method for attempting to access a first data entity in a file system, the method being performed by a first computing entity, the file system also including one or more additional data entities that are concurrently accessible to at least one other computing entity, the method comprising:

the first computing entity attempting to obtain a lease for itself on the first data entity without needing another computing entity acting on its behalf, by reading, using a processor, an owner field included in the file system that can be used to determine whether the first data entity is leased by a computing entity and a time field included in the file system that can be used to determine whether a lease for the first data entity has expired and:

if the owner field indicates that the first data entity is not currently leased, the first computing entity writing to the owner field in the file system to indicate an assumption of a lease of the first data entity and writing to the time field in the file system to indicate when the lease expires; else ~~or~~

if the owner field indicates that the first data entity has been leased, the first computing entity reading the time field in the file system and:

if the time field indicates that the lease has expired, the first computing entity writing to the owner field in the file system to break the existing lease and to indicate an assumption of a new lease and the first computing entity writing to the time field in the file system to indicate when the new lease expires; else ~~or~~

if the time field indicates that the lease is still active, concluding that the first data entity is currently unavailable; and

if a lease is obtained, the first computing entity accessing the first data entity while the lease is in effect.

2. (Previously presented): The method of claim 1, wherein the first data entity is a file.

3. (Previously presented): The method of claim 2, wherein the first data entity includes metadata and the owner field is located in this metadata.

4. (Previously presented): The method of claim 1, wherein the first data entity is a directory.

5. (Previously presented): The method of claim 1, wherein the step of the first computing entity writing to the owner field to indicate an assumption of a lease of the first data entity comprises the first computing entity writing a data value to the owner field that uniquely identifies the first computing entity.

6. (Previously presented): The method of claim 5, wherein the data value that uniquely identifies the first computing entity is determined autonomously by the first computing entity.

7. (Previously presented): The method of claim 5, wherein the owner field indicates that the first data entity is not currently leased when the owner field contains a value of zero.

8. (Previously presented): The method of claim 1, wherein a lease expires a predetermined period of time after the lease begins, and wherein the step of the first computing entity writing to the time field to indicate when the lease expires comprises the first computing entity writing a current time value to the time field.

9. (Previously presented): The method of claim 1, wherein the first computing entity determines whether a prior lease has expired by reading a first value from the time field, delaying for a predetermined lease period and reading a second value from the time field, wherein the first computing entity determines that the prior lease has expired if the second value is the same as the first value, and the first computing entity determines that the prior lease has not expired if the second value is different from the first value.

10. (Previously presented): The method of claim 1, wherein the steps of reading the owner field and reading the time field are both performed in a single read operation.

11. (Previously presented): The method of claim 1, wherein, if the first computing entity concludes that the first data entity is currently unavailable, the first computing entity further writes an entry to a queue owner field in a queue in the file system to indicate an interest in accessing the first data entity.

12. (Previously presented): The method of claim 11, wherein the first computing entity also writes to a queue time field in the queue in the file system to indicate a period of time for which the entry to the queue owner field is valid.

13. (Previously presented): The method of claim 1 further comprising the first computing entity reserving a disk on which the owner field and the time field are located to ensure exclusive access to the disk for the reading and writing of the owner field and the time field.

14. (Previously presented): The method of claim 1, wherein, if a lease is obtained, the first computing entity also sets a renewal timer and, after the renewal timer expires, the first computing entity renews the lease by writing a new value to the time field.

15. (Withdrawn): A computer system comprising a first physical computer, a second physical computer, a data storage unit, a first data link for connecting the first physical computer to the data storage unit and a second data link for connecting the second physical computer to the data storage unit, the computer system further comprising:

a first virtual machine running on the first physical computer;

a second virtual machine running on the second physical computer; and

a file system stored on the data storage unit, the file system comprising:

a first data entity, the first data entity being usable by the first virtual machine and by the second virtual machine; and

a lock for providing exclusive access to the first data entity, the lock comprising an owner field and a time field, the owner field being used to determine if the first data entity has been leased by a computing entity and the time field being used to determine when a lease of the first data entity expires.

16. (Withdrawn): The computer system of claim 15, wherein the first data entity is a file.

17. (Withdrawn): The computer system of claim 16, wherein the file system further comprises a second file implementing a first virtual disk drive for use by the first virtual machine and a third file implementing a second virtual disk drive for use by the second virtual machine.

18. (Withdrawn): The computer system of claim 17, wherein the first data entity is the second file.

19. (Withdrawn): The computer system of claim 15, wherein the first data entity is a directory.

20. (Withdrawn): The computer system of claim 15, wherein the file system further comprises a queue that may be used by a computing entity to indicate an interest in accessing the first data entity in the event that another computing entity has exclusive access to the first data entity.

21. (Withdrawn): The computer system of claim 15, wherein the first data entity includes metadata and the lock is located in this metadata.

22. (Withdrawn): The computer system of claim 15, wherein the first data link and the second data link are part of a data storage network.

23. (Withdrawn): The computer system of claim 15, wherein the data storage unit comprises a disk drive.

24. (Withdrawn): The computer system of claim 23 further comprising a disk reservation capability for providing exclusive access to the disk when accessing the lock.

25. (Withdrawn): The computer system of claim 15 wherein the first physical computer autonomously determines a first unique data value for identifying the first virtual machine in the owner field and the second physical computer autonomously determines a second unique data value for identifying the second virtual machine in the owner field.

26. (Withdrawn): The computer system of claim 15 wherein the first virtual machine is migrated from the first physical computer to the second physical computer and the first data entity remains usable by the first virtual machine and by the second virtual machine.

27. (Withdrawn): The computer system of claim 26 wherein the first data entity is a primary virtual disk drive for use by the first virtual machine.

28. (Withdrawn): The computer system of claim 15 wherein, if the first data entity has been leased for use by the first virtual machine and the first physical computer fails so that the lease on the first data entity cannot be released by the first physical computer, when the lease on the first data entity expires, the second physical computer is able to break the lease on the first data entity and begin using the first data entity.

29. (Withdrawn): The computer system of claim 28 wherein the first data entity is a virtual disk drive for the first virtual machine and, after the failure of the first physical computer and after the breaking of the lease on the first data entity, the first virtual machine is restarted on the second physical computer using the first data entity.

30. (Previously presented): A method for attempting to access a first data entity in a file system, the method being performed by a first computing entity, the file system also including one or more additional data entities that are concurrently accessible to at least one other computing entity, the method comprising:

the first computing entity attempting to access the first data entity without needing another computing entity acting on its behalf, reading, using a processor, an owner field included in the file system that can be used to determine whether the first data entity is in use by a computing entity and determining whether the first data entity is in use by a computing entity;

if the first data entity is not in use by a computing entity, the first computing entity writing to the owner field in the file system to take control of a lock on the first data entity; and

if control of the lock is obtained, the first computing entity accessing the first data entity; else ~~or~~

if control of the lock is not obtained, the first computing entity writing an entry to a queue owner field in the file system to indicate an interest in accessing the first data entity and waiting for an opportunity to access the first data entity.

31. (Previously presented): The method of claim 30 further comprising, if the first data entity is in use by a computing entity, the first computing entity reading a time field in the file system to determine whether a lease on the data entity has expired and, if the lease has expired, the first computing entity writing to the owner field to break the existing lease and to indicate an assumption of a new lease of the first data entity.

32. (Previously presented): The method of claim 31, wherein the first computing entity determines whether the lease has expired by reading a first value from the time field, delaying for a predetermined lease period and reading a second value from the time field, wherein the first computing entity determines that the lease has expired if the second value is the same as the first value, and the first computing entity determines that the lease has not expired if the second value is different from the first value.

33. (Previously presented): The method of claim 30 further comprising, if the first data entity is not in use by a computing entity, in addition to writing to the owner field to take control of the lock on the first data entity, the first computing entity writing to a time field in the file system to indicate when a lease of the first data entity expires.

34. (Previously presented): The method of claim 30, wherein the first data entity is a file.

35. (Previously presented): The method of claim 34, wherein the first data entity includes metadata and the owner field is located in this metadata.

36. (Previously presented): The method of claim 30, wherein the first data entity is a directory.

37. (Previously presented): The method of claim 30 further comprising the first computing entity reserving a disk on which the owner field is located to ensure exclusive access to the disk for the reading and writing of the owner field.

38. (Previously presented): The method of claim 30, wherein the first computing entity autonomously determines a data value that uniquely identifies the first computing entity and the first computing entity assumes a lock on the first data entity by writing the unique data value into the owner field.

39. (Previously presented): The method of claim 30 further comprising, if control of the lock is not obtained, in addition to writing an entry to a queue owner field in the file system to indicate an interest in accessing the first data entity, the first computing entity writing to a queue time field in the file system to indicate a period of time for which the entry to the queue owner field is valid.

40. (New): The method of claim 1 wherein the owner field is stored in a file system lock structure in the file system.